# Linux Virtualisation

Stephen Quinney
<squinney@inf.ed.ac.uk>

# Virtualisation is...

*“... a framework or methodology of dividing the resources of a computer into multiple execution environments, by applying one or more concepts or technologies such as hardware and software partitioning, time-sharing, partial or complete machine simulation, emulation, quality of service, and many others.”*

# Reasons to Virtualise

- Consolidation
- Containment
- Management
- Legacy support
- Quality of service
- Testing and debugging

# Options

- ABI/API Emulation (e.g Wine)
- Full hardware emulation (e.g qemu)
- Containers/Zones - “jail”
- Programming language VMs
- Binary translation
- Para-virtualisation (UML, Xen)
- Hardware-assisted virtualisation

# Process Virtualisation

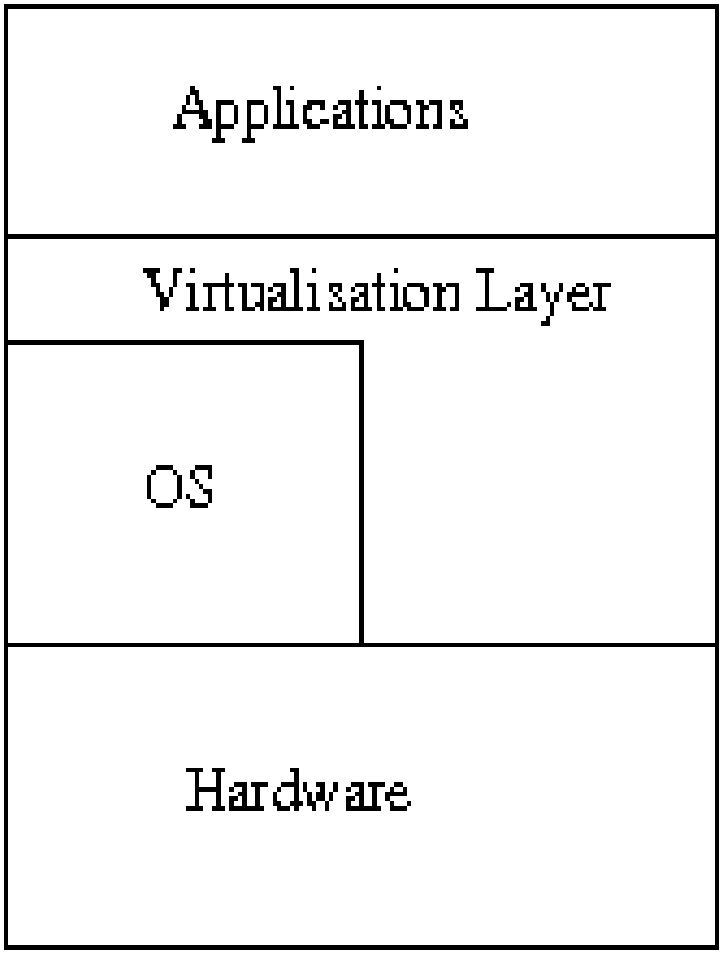

# System Virtualisation

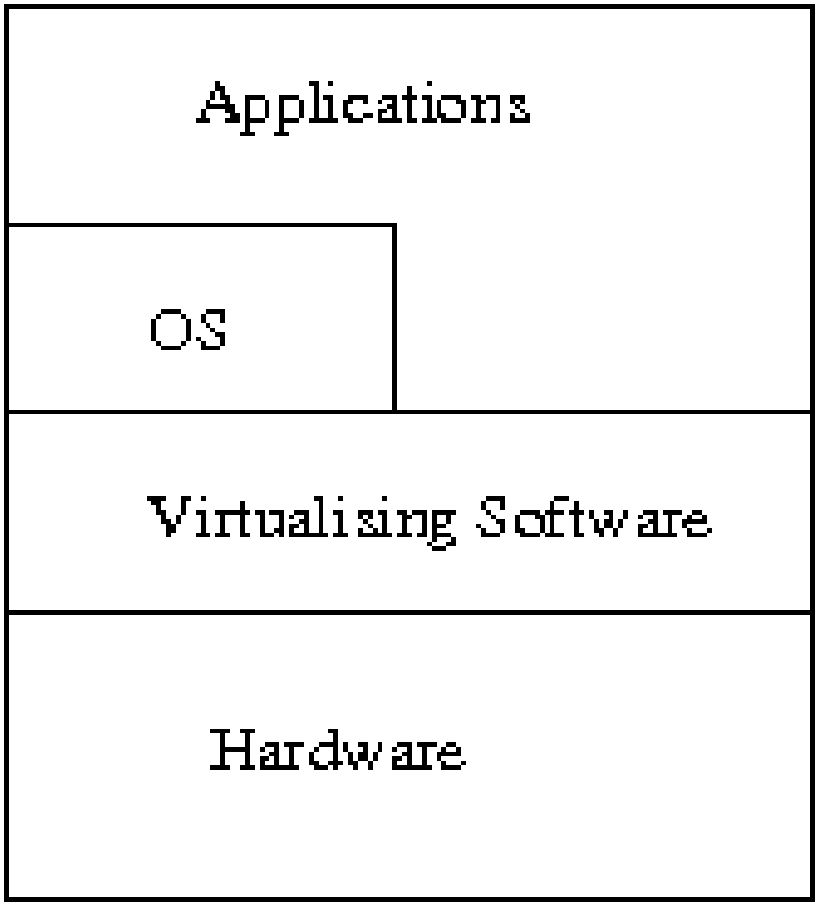

# System Virtualisation

- Has been around for a long time.
- There is a *Virtual Machine Monitor*, sometimes known as a *hypervisor.*
- There are a set of guest systems.

# Requirements

- **Equivalence:** Software running in a VM must act the same as on a “real” machine
- **Performance:** Majority of instructions should be executed natively without intervention from the hypervisor
- **Resource control:** hypervisor must have complete control of hardware resources

# IA-32 Problems

- Never designed for virtualisation
- Not all instructions that should “trap” actually do
- Some instructions act differently in different privilege modes
- Some privileged machine state (e.g. page tables) is visible to user-mode software.

# Para-virtualisation

- Modify the guest operating system to make it aware that it is virtualised
  - Xen
  - VMware
- Very fast – almost native speed

# Binary Rewriting

- The hypervisor rewrites binaries in memory to avoid problematic instructions
  - Vmware ESX Server
- Allows running of any unmodified OS which is the same architecture.
- Not as fast as paravirtualisation.

# Hardware Virtualisation

- Allows running of any unmodified OS as a virtualised guest.
- New processors have extensions to support virtualisation
  - Intel VT (*vmx*)
  - AMD-SVM (*svm)*
- Supported by Xen, Linux KVM

# Some nice features

- Complete separation/containment
- Live migration
- Resource control